



AF L
IFW

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Application No.: 09/697,730
Filed: October 25, 2000
Inventor(s):
Cassiday et al.

Examiner: Sefcheck, Gregory
Group/Art Unit: 2662
Atty. Dkt. No: 5681-96500

Title: HIGH PERFORMANCE
TRANSMISSION LINK
AND INTERCONNECT

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on the date indicated below.

Stephen J. Curran

Printed Name



Signature

9/2/04

Date

APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir/Madam:

Further to the Notice of Appeal of July 15, 2004, Appellants present this Appeal Brief. Appellants respectfully request that this appeal be considered by the Board of Patent Appeals and Interferences.

09/08/2004 BSAYAS11 00000020 501505 09697730
01 FC:1402 330.00 DA

I. REAL PARTY IN INTEREST

The subject application is owned by Sun Microsystems, Inc., a corporation organized and existing under and by virtue of the laws of the State of Delaware, and having its principal place of business at 4150 Network Circle, Santa Clara, CA 95054, as evidenced by the assignment recorded at Reel 011755, Frame 0608.

II. RELATED APPEALS AND INTERFERENCES

No other appeals or interferences are known which would directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Claims 18-28 and 31-34 are pending in the present application. Claims 18-28 and 31-34 stand finally rejected and are the subject of this appeal. A clean copy of claims 18-28 and 31-34, as on appeal (incorporating all amendments), is included in the Appendix hereto.

IV. STATUS OF AMENDMEMNTS

No amendment to the claims has been filed subsequent to the final rejection. The Appendix hereto reflects the current state of the rejected claims.

V. SUMMARY OF THE INVENTION

In a typical computer network, as users are added and applications become more complex, available bandwidth is quickly consumed. It is important that computer networks operate and be managed as efficiently as possible. For example, one aspect of increasing overall performance of a network is to increase the throughput of network nodes by increasing the throughput of packets from receivers to transmitters that use a crossbars or switches in a given node.

However, in conventional node receivers, when a packet is received at a node on a particular communication link, it is placed in a single receiver buffer that is associated with that link. In such receivers, one problem that is packets wanting to go to different transmitters may be blocked in the single buffer until the packet that is ahead of the waiting packets are routed first. This is the typical nature of a queue. For example, a first packet needs to go out to transmitter 0 and another packet behind it in a queue in the same buffer needs to go to transmitter 1. In one clock cycle, the packet for Xmit 0 gets out, and in the next clock cycle, the packet behind it for Xmit 1 gets out. However, the second packet had to wait one clock cycle before it could be sent out to a different transmitter which may have been idle waiting for it. In other words, the second packet was blocked from getting out because the receiver has only one buffer. This blocking problem may become exponentially worse as the number of receivers and transmitters increases since it is more likely that packets in a receiver buffer will be destined for different transmitters. (See specification pages 24 and 25)

Thus, it may be desirable to improve the throughput in a node. According to one embodiment of Appellant's claimed invention, a receiver in the node has two buffers. When the receiver gets a data segment, such as 8 bits of a flit, it examines a particular bit and determines, based on the bit, whether the data segment should go to a first buffer or a second buffer. For example, if the bit is a one, it will go to the first buffer and if it is a zero it will go to the second buffer. The node also has a dual crossbar, one for each buffer. A first crossbar receiving data segments from one buffer and the second crossbar receives data segments from the other buffer. This allows the potential routing of two data segments in one clock cycle to their respective transmitters so that one segment does not have to wait for the other segment to be routed.

In another aspect of the Appellant's invention, a method of routing a received data packet through a node is described. A data packet is received at a receiver in the node. The packet is examined based on one or more categorical bits in the data packet, such as a

stripe bit. The data packet is then sorted based on the categorical bits and sent to one of multiple buffers. The packet is then inputted to a crossbar that corresponds to the buffer from which the packet came. The packet is then routed to a transmitter such that two data packets can be processed by the node in one clock cycle. In one embodiment, the order of sequential data packets passing through one of the buffers and crossbars is maintained so that packets that belong together and have a certain order are one of the plurality of buffers. (See, e.g., FIG. 9 and specification pages 5 and 6.

As discussed for one embodiment at page 25, lines 5-22 “Shown in FIG. 9 are three receivers, RCV0, RCV1, and RCV2. Each receiver has two buffers for holding the incoming data traffic, which is split into each buffer as described below. RCV0 has two buffers, buffer1 902 and buffer2 904, and a buffer select logic unit 906 which essentially splits the incoming traffic to one of the two buffers based on a predetermined criteria. The other components normally found in a receiver but not relevant to the dual crossbar concept are not shown in FIG. 9. Each of the other two or more receivers also have two receiver buffers, shown as buffers 908 and 910 in RCV 1 and buffers 912 and 914 in RCV 2.

Each packet has information that can be used to categorize a packet that is used by buffer select logic 906. For example, an address, such as an address for Node 2 or Node 1, contained in the packet can be used to categorize the packet. In the described embodiment, it is important to maintain order over a set of addresses or on an address-to-address basis. As mentioned, one way of increasing the efficiency of routing packets through a node is to have a second receiver buffer (e.g., buffers 904, 910 and 914) for each receiver and a second crossbar in the node.”

As discussed for one embodiment at page 27, lines 4-13 “A transmitter can receive packets from either crossbar so long as the order of packets within a stripe is maintained. That is, all packets with a stripe bit of 1, for example, are in the correct order when being routed through the crossbars and to the transmitters. A transmitter has two input buffers, for example, buffer1 920 and buffer2 922. Buffer1 920 (as well as the

other buffer1s in XMIT1 and XMIT2) receives input from crossbar 916 (the A MUXs) and buffer2 922 receives input from crossbar 918 (as well as the other buffer2s in XMIT1 and XMIT2). When the transmitter gets packets from crossbar 916 and crossbar 918, it decides using an arbitrator 924 which packet will be sent out on the link.”

VI. ISSUES

1. Whether claims 18, 20-23, 25, and 27-28 are anticipated by Suzuki (U.S. Patent Number 6,625,160) under 35 U.S.C. § 102(e).

2. Whether claims 19, 24, and 26 are patentable over Suzuki in view of Kessler (U.S. Patent Number 6,567,900) under 35 U.S.C. § 103(a).

3. Whether claims 31-34 are patentable over Suzuki in view of McGill (U.S. Patent Number 5,436,886) under 35 U.S.C. § 103(a).

VII. GROUPING OF CLAIMS

Claims 18-22 stand or fall together.

Claims 23-28 stand or fall together.

Claims 31-34 stand or fall together.

The above claim groupings are for purposes of this appeal only. The reasons why each group of claims is believed to be separately patentable are explained below in the Argument.

VIII. ARGUMENT

A. Claims 18, 20-22

The Examiner rejected claims 18 and claims 20-22 as being anticipated by Suzuki under 35 U.S.C. § 102(e). Appellants respectfully traverse this rejection in light of the following remarks.

The Examiner also rejected claim 19 as being obvious over Suzuki in view of Kessler under 35 U.S.C. § 103(a). Appellants respectfully traverse this rejection in light of the following remarks.

Suzuki **does not teach or disclose** each and every feature recited in Appellant's claim 18. Specifically, Suzuki **does not teach or disclose** "a first crossbar for receiving the first data segment from the first buffer; and a second crossbar for receiving the second data segment from the second buffer, such that the first data segment and the second data segment are routed to one or more transmitters in one clock cycle in the node," as recited in Appellant's claim 18.

Suzuki is directed to a minimum bandwidth guarantee for a cross-point buffer switch, wherein Suzuki discloses at col. 5, lines 59-65

"It has been discovered that a method and system can be achieved which will substantially alleviate the starving of relatively low or lower priority data packets at a particular crossbar switch within a data communications network while substantially ensuring that relatively high or higher priority data packets transit that same particular crossbar switch with relatively low delay."

In addition, Suzuki discloses at col. 9, lines 5-7 "Examples of queues are illustrated in FIG. 5 as ingress-egress queues 516-526, each of which is shown respectively connected between its associated particular ingress-egress pair."

Suzuki further discloses at col. 10, lines 57-64

"Illustrated is that both switch egress_1 512 and switch egress_2 514 have respectively associated with them (1) excess bandwidth schedulers 600, 604, and (2) guaranteed minimum bandwidth schedulers 602, 606. Each guaranteed minimum bandwidth scheduler for each egress assures that each queue, and consequently each ingress associated with each queue, is

given some "guaranteed," or target, minimum data transmission per unit time"

From the foregoing, it appears to Appellant that Suzuki is teaching a system that prevents starvation of low priority packets by higher priority packets, wherein each of the buffers 516-524 are dedicated buffers associated with specific egress ports. Accordingly, the schedulers 528/530 are not crossbars, but schedulers that determine which ingress/egress pair is selected based on one or more minimum guaranteed bandwidth schemes. The Applicant submits that based upon the drawings of FIG. 5 and FIG. 6, it appears that switch egress 512 and 514 and schedulers 600-606 use two, (three input to one output) switches to select an ingress/egress pair.

Accordingly, Applicant believes that claim 18, along with its dependent claims, patentably distinguishes over Suzuki and over Suzuki in view of Kessler for the reasons given above.

Furthermore, notwithstanding the above, in his rejection of claim 31 (discussed below), the Examiner acknowledges that **Suzuki does not show a first or second crossbar** that connects data from the first and second buffer to any of the plurality of output lines. Accordingly, the applicant respectfully submits that the Examiner's rejection of claims 18, 20-23, 25, and 27-28 under 35 U.S.C. §102(e) is improper. The Applicant requests that the rejection be withdrawn.

B. Claims 23-28

The Examiner rejected claim 23, 25, and 27-28 as being anticipated by Suzuki under 35 U.S.C. § 102(e). Appellants respectfully traverse this rejection in light of the following remarks.

The Examiner also rejected claim 24, and 26 as being obvious over Suzuki in view of Kessler under 35 U.S.C. § 103(a). Appellants respectfully traverse this rejection in light of the following remarks.

Claims 23-28 are distinguishable for the reasons given above in regard to claim 18. In addition, Suzuki **does not teach or disclose** “sorting the data packet to one of a plurality of buffers based on the one or more categorical bits in the data packet; and inputting the data packet to one or more crossbars, a crossbar corresponding to a buffer, and routing the data packet to a transmitter such that two data packets can be processed by the node in one clock cycle,” as recited in Appellant’s claim 23.

Accordingly, Appellant’s submit that claim 23, along with its dependent claims 24-28, patentably distinguishes over Suzuki and over Suzuki in view of Kessler for the reasons given above.

C. Claim 31-34

The Examiner rejected claims 31-34 under 35 U.S.C. § 103(a) as being unpatentable over Suzuki in view of McGill. Appellants respectfully further traverse this rejection in light of the following remarks.

The Examiner has acknowledged that **Suzuki does not show a first or second crossbar that connects data from the first and second buffer to any of the plurality of output lines**. However, the Examiner asserts that McGill discloses an ATM switch for routing packets between multiple input and output ports. The Examiner further asserts that it would have been obvious “to modify the switching node of Suzuki by enabling the first and second crossbars to connect data from the first and the second buffers to any of the plurality of output lines.” The Examiner’s rationale for such a modification is stated as “such a modification would provide redundancy to enable that data from either buffer may be connected to its appropriate output line if one of the two crossbars failed.” The

Applicant respectfully disagrees with both the Examiner's assertion that it would have been obvious and his motivation to do so.

Specifically, the McGill reference is directed to an ATM switch of dual plane operation (i.e., the dual synchronous matched scheme (dual plane operation)) and which withstands multiple faults for exchanging cells among a plurality of bidirectional ports through a first and a second switch plane. Each of the plurality of bidirectional ports is connected to a plurality of linecards to receive the cells therefrom and transmit the cells thereto. The ATM switch comprises a first and a second switch fabric connected to the plurality of bidirectional ports for transferring the cells among bidirectional ports. (See McGill col. 2, lines 12-13 and 51-60)

McGill also discloses at col. 3 line 63 through col. 4 line 2

“The switch is fully redundant in that each port has redundant multiplexers (AX) 14 denoted AX0 and AX1 and the switch fabric 16 is also redundant, consisting of switch fabrics labeled SF0 in one switch plane and SF1 in another. Any number of linecards 18 can be attached to each port but must have an aggregate bandwidth of 600 Mb/s. Each linecard is also protected by pair (1+1) redundancy. The figure shows one protected pair of linecards attached to each port and the pair are labeled LC0 and LC1. The connection between linecards and multiplexers is full crossover on ingress and egress, that is to say, it is a point-to-2-point in both directions.”

From the foregoing, it appears to the Appellant that McGill teaches a fully redundant multiply-fault tolerant system which may provide great utility in systems where faults may be common and also detrimental. However, from the discussion above, Suzuki is directed to providing minimum bandwidth guarantees in a switch, not fault tolerance. Appellant can find no reference in Suzuki to failures or redundancy in the event of a failure. Thus, Suzuki is not directed toward solving the same or even a similar problem as McGill.

Appellant further submits that since the ingress buffers in Suzuki are dedicated to specific egress ports, there is absolutely no motivation to combine the references. The

Appellant further doubts that any person skilled in the art would be motivated to combine a system of McGill's complexity with Suzuki as suggested by the Examiner.


The MPEP states at § 2143.01 "The fact that the references can be combined or modified is not sufficient to establish *prima facie* obviousness. The mere fact that the references can be combined or modified does not render the resultant combination obvious unless the prior art also suggests the desirability of the combination. *In re Mills*, 916 F.2d 680, 16 USPQ2d 1430 (Fed. Cir. 1990)... Although a prior art device "may be capable of being modified to run the way the apparatus is claimed, there must be a suggestion or motivation in the reference to do so." 916 F.2d at 682, 16 USPQ2d at 1432.). See also *In re Fritch*, 972 F.2d 1260, 23 USPQ2d 178 (Fed. Cir. 1992)..."

Accordingly, Applicant believes that claim 31, along with its dependent claims 32 through 34, patentably distinguishes over Suzuki in view of McGill for the reasons given above.

IX. CONCLUSION

For the foregoing reasons, it is submitted that the Examiner's rejection of claims 18-28 and 31-34 was erroneous, and reversal of his decision is respectfully requested.

Respectfully submitted,



Stephen J. Curran

Reg. No. 50,664

AGENT FOR APPLICANT(S)

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.

P.O. Box 398

Austin, TX 78767-0398

Phone: (512) 853-8800

Date: 9/2/04

X. APPENDIX

The claims on appeal are as follows.

18. (Original) A node in an interconnect link system comprising:
a first buffer for receiving a first data segment passing a first criteria based on a predetermined one or more bits for the first segment;
a second buffer for receiving a second data segment passing a second criteria based on the predetermined one or more bits for the second segment;
a first crossbar for receiving the first data segment from the first buffer; and
a second crossbar for receiving the second data segment from the second buffer, such that the first data segment and the second data segment are routed to one or more transmitters in one clock cycle in the node.
19. (Previously presented) A node as recited in claim 18 further comprising a data packet having a plurality of bits, the predetermined one or more bits being a stripe bit wherein the stripe bit is used for determining the appropriate buffer to sort data segments into.
20. (Original) A node as recited in claim 18 further comprising a receiver capable of sorting a plurality of received data segments based on the predetermined one or more bits in a data segment.
21. (Original) A node as recited in claim 18 further comprising a transmitter having an arbitrator to decide which data segment to transmit.
22. (Original) A node as recited in claim 18 wherein the first buffer and the second buffer are in a receiver.
23. (Original) A method of routing a received data packet through a node, the method comprising:

receiving a data packet at a receiver in the node;
examining the data packet based on one or more categorical bits in the data packet;
sorting the data packet to one of a plurality of buffers based on the one or more categorical bits in the data packet; and
inputting the data packet to one or more crossbars, a crossbar corresponding to a buffer, and routing the data packet to a transmitter such that two data packets can be processed by the node in one clock cycle.

24. (Original) A method as recited in claim 23 wherein examining the data packet further includes determining whether a stripe bit in the data packet is zero or one.

25. (Original) A method as recited in claim 23 wherein sorting the data packet further includes routing the data packet to a first buffer if the one or more categorical bits meets a first criteria and routing the data packet to a second buffer if the one or more categorical bits meets a second criteria.

26. (Original) A method as recited in claim 25 wherein the first criteria is that one or more of the categorical bits be a zero and the second criteria is that one or more of the categorical bits be a one.

27. (Original) A method as recited in claim 23 wherein inputting the data packet to one or more crossbars further comprises routing the data packet to a transmitter.

28. (Original) A method as recited in claim 23 further comprising maintaining the order of sequential data packets passing through one of the plurality of buffers.

31. (Previously presented) A routing node suitable for use in a network that carries data packets, the routing node having a plurality of input lines and a plurality of output lines, the node comprising:

a first receiver that receives packets from a first input line, the first receiver including a first buffer arranged to receive data packets that contain one or more selected bits that meet a first predetermined criteria, and a second buffer arranged to receive at least some of the data packets that are not directed to the first buffer;

a first crossbar arranged to connect data packets from the first buffer with any of a plurality of output lines; and

a second crossbar arranged to connect data packets from the second buffer with any of the plurality of output lines, whereby packets received by the first input line may be transmitted to an appropriate output line through either the first or second crossbar.

32. (Previously presented) A routing node as recited in claim 31 comprising a plurality of receivers, each receiver being arranged to receive packets from an associated input line, and wherein each receiver has associated first and second buffers, each first buffer being coupled to the first crossbar and each second buffer being coupled to the second crossbar.

33. (Previously presented) A routing node as recited in claim 32 wherein each receiver is arranged to peek at a designated stripe bit in each data packet received by the receiver, wherein if the stripe bit is a designated value, the data packet is passed to the first data buffer.

34. (Previously presented) A routing node as recited in claim 33 wherein if the stripe bit is not the designated value, the data packet is passed to the second data buffer.



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Application No.: 09/697,730
Filed: October 25, 2000
Inventor(s):
Cassiday et al.

Title: HIGH PERFORMANCE
TRANSMISSION LINK
AND INTERCONNECT

§ Examiner: Sefcheck, Gregory
§ Group/Art Unit: 2662
§ Atty. Dkt. No: 5681-96500
§
§
§
§
§
§
§
§
§
§

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on the date indicated below.

Stephen J. Curran

Printed Name



Signature

September 2, 2004

Date

FEE AUTHORIZATION


Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

The Commissioner is hereby authorized to charge the following fee to Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C. Deposit Account Number 50-1505/5681-96500:

Fee: Appeal Brief
Amount: \$330.00
Attorney Docket No.: 5681-96500

The Commissioner is also authorized to charge any extension fee or other fees which may be necessary to the same account number.

Respectfully submitted,


Stephen J. Curran
Reg. No. 50,664
AGENT FOR APPLICANT(S)

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8800
Date: September 2, 2004